

PICO: Mitigating Heterodyne Crosstalk Due to Process Variations and Intermodulation Effects in Photonic NoCs

Sai Vineel Reddy Chittamuru, Ishan G Thakkar, Sudeep Pasricha
Department of Electrical and Computer Engineering
Colorado State University, Fort Collins, CO, U.S.A.
{sai.chittamuru, ishan.thakkar, sudeep}@colostate.edu

ABSTRACT

Photonic networks-on-chip (PNoCs) employ photonic waveguides with dense-wavelength-division-multiplexing (DWDM) for signal traversal and microring resonators (MRs) for signal modulation, to enable high bandwidth on-chip transfers. Unfortunately, DWDM increases susceptibility to intermodulation effects, which reduces signal-to-noise ratio (SNR) for photonic data transfers. Additionally, process variations induce variations in the width and thickness of MRs causing resonance wavelength shifts, which further reduces SNR, and creates communication errors. This paper proposes a novel framework (called *PICO*) for mitigating heterodyne crosstalk due to process variations and intermodulation effects in PNoC architectures. Experimental results indicate that our approach can improve the worst-case SNR by up to 4.4× and significantly enhance the reliability of DWDM-based PNoC architectures.

Categories and Subject Descriptors: C.4 [Performance of Systems]: Reliability, availability, and serviceability

General Terms – Reliability, Performance, Experimentation

Keywords – Process Variations, Crosstalk, Photonic NoCs

1. INTRODUCTION

Recent developments in the area of silicon photonics have enabled the integration of photonic components with CMOS circuits on a chip. Several photonic network-on-chip (PNoC) architectures have been proposed to date (e.g., [1]-[3]). These architectures employ on-chip photonic links that use microring resonator (MR) modulators to modulate electrical signals onto photonic signals that travel through a silicon waveguide, and MR filter receivers that detect and drop photonic signals on to a photodetector to recover an electrical signal. Each MR has a unique set of (resonance) wavelengths that it can couple to and work correctly with. Typically, silicon waveguides are designed to support dense wavelength division multiplexing (DWDM), where a large number of wavelengths are multiplexed in the waveguide. The use of multiple MRs that are in resonance with these wavelengths enables high bandwidth parallel data transfers.

Prior work indicates that heterodyne crosstalk is a major contributor of crosstalk noise in DWDM-based waveguides, which reduces photonic signal SNR and reliability in PNoCs [4]. Heterodyne crosstalk noise occurs at a detector MR when it picks up some non-resonant optical power from neighboring wavelengths. The strength of the heterodyne crosstalk noise at a detector MR depends on the following three attributes: (i) channel gap between the MR resonant

wavelength and the adjacent wavelengths; (ii) Q-factors of neighboring detector MRs, and (iii) the strengths of the non-resonant signals at the detector. With increase in DWDM, the channel gap between two adjacent wavelengths decreases, which in turn increases heterodyne crosstalk in detector MRs. With decrease in Q-factors of MRs, the widths of the resonant passbands of MRs increases, increasing passband overlap among neighboring MRs, which in turn increases heterodyne crosstalk. The strengths of the non-resonant signals depend on the losses faced by the non-resonant signals throughout their path from the laser source to the MR detector.

Intermodulation (IM) crosstalk has the biggest influence on the last attribute discussed above, causing loss of non-resonant signals in a DWDM waveguide [5]. IM crosstalk occurs when a modulator MR truncates and consequently modulates the passbands of the neighboring non-resonant signals. Thus the level of heterodyne crosstalk and resultant SNR at the detector depends on the amount of IM passband truncation at the modulator. This motivates mitigating the effects of IM passband truncation on heterodyne crosstalk by controlling the strengths of the non-resonant signals at the detector.

Additionally, fabrication process variations (PV) induce variations in the width and thickness of MRs, which cause resonance wavelength shifts in MRs [6][8]. PV-induced resonance shifts reduce the channel gap between the resonances of the victim MRs and adjacent MRs, which increases crosstalk and worsens SNR. The worsening of SNR deteriorates the bit-error-rate (BER) in a waveguide. For example, a previous study shows that in a DWDM-based photonic interconnect, when PV-induced resonance shift is over 1/3 of the channel gap, BER increases from 10^{-12} to 10^{-6} [9]. Techniques to counteract the PV-induced resonance shifts in MRs involve realigning the resonant wavelengths by using localized trimming [8] or thermal tuning [6]. Localized trimming is the more viable technique as it enables faster and finer grained control that is also not impacted by on-die thermal variations, unlike thermal tuning. However, our analysis has shown that localized trimming increases intrinsic optical signal loss in MRs and waveguides due to the free carrier absorption effect (FCA). This loss decreases Q-factor of MRs, which increases heterodyne crosstalk in MRs and reduces SNR.

In this paper, we present a novel crosstalk mitigation framework called *PICO* to enable reliable communication in emerging PNoC-based multicore systems. *PICO* mitigates the effects of IM crosstalk by controlling signal loss of wavelengths in the waveguide and reduces trimming-induced crosstalk by intelligently reducing undesirable data value occurrences in a photonic waveguide based on the PV profile of MRs. Our framework has low overhead and is easily implementable in any existing DWDM-based PNoC without major modifications to the architecture. To the best of our knowledge, this is the first work that attempts to improve SNR in PNoCs considering both IM effects and PV in its MRs. Our novel contributions are:

- We present device-level analytical models to capture the deleterious effects of localized trimming in MRs. Moreover, we extend this model for system-level heterodyne crosstalk analysis;
- We propose a scheme for IM passband truncation-aware heterodyne crosstalk mitigation (IMCM) to improve worst-case

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DAC '16, June 05-09, 2016, Austin, TX, USA

© 2016 ACM. ISBN 978-1-4503-4236-0/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2897937.2898063>

SNR of MRs by controlling non-resonant signal power;

- We propose a scheme for PV-aware heterodyne crosstalk mitigation (PVCMM) to improve worst-case SNR of detector MRs by encoding data to avoid undesirable data occurrences;
- We evaluate our proposed *PICO* (PVCMM+ IMCM) framework by implementing it on the well-known Corona crossbar PNoC architecture [1], [10], and compare it with two encoding based heterodyne crosstalk mitigation mechanisms from [12] for real-world multi-threaded PARSEC benchmarks.

2. BACKGROUND AND RELATED WORK

DWDM-based PNoCs utilize several photonic devices such as microring resonators (MRs) as modulators and detectors, photonic waveguides, splitters, and trans-impedance amplifiers (TIAs). The reader is directed to [12] for more discussion on these devices.

An important characteristic of photonic signal transmission in on-chip waveguides is that it is inherently *lossy*, i.e., the light signal is subject to losses such as through-loss in MR modulators and detectors, modulating losses in modulator MRs, detection loss in detector MRs, propagation and bending loss in waveguides, and splitting loss in splitters. Such losses negatively impact SNR in waveguides. In addition to the optical signal loss, crosstalk noise of the constituent MRs also deteriorates SNR. Both modulators and detectors are susceptible to crosstalk noise in DWDM-based PNoCs.

Crosstalk noise can be classified as *homodyne* or *heterodyne*. Homodyne crosstalk usually occurs in MRs used as optical injectors, when an injector MR couples optical power of the same wavelength from two different ports to a single output port. Heterodyne crosstalk occurs in detector and modulator MRs when an MR picks up some optical power from non-resonant signals. As discussed in [4], homodyne crosstalk may either contribute to the noise or cause fluctuation in the signal power, which makes the analysis and mitigation of homodyne crosstalk more complicated and beyond the scope of this work. Thus this work focuses on heterodyne crosstalk and propose solutions to mitigate it. In the rest of the paper, we use the term crosstalk to refer heterodyne crosstalk.

A few prior works have analyzed crosstalk in PNoCs. The effect of crosstalk noise on SNR is shown to be negligible in WDM systems presented in [13] and [14], as these systems use only four WDM wavelengths per waveguide. In [5], IM effects are shown to be negligible for a WDM link operating at 10 Gb/s. However, in PNoC architectures that use DWDM (e.g., Corona [1] with 64 wavelength DWDM), there exists significant crosstalk noise. The damaging impact of crosstalk noise in the Corona PNoC is presented in [15], where worst-case SNR is estimated to be 14dB in data waveguides, which is insufficient for reliable data transfers. To mitigate the impact of crosstalk noise in DWDM based PNoCs, two encoding techniques (PCTM5B and PCTM6B) were presented in [12]. In [7] a technique was proposed to increase channel spacing between adjacent DWDM wavelengths, to mitigate crosstalk in MR detectors. However, *none of these works considers the system-level impact of IM effects or PV on crosstalk in DWDM-based PNoCs.*

Fabrication-induced process variations (PV) impact the cross-section, i.e., width and height, of photonic devices such as MRs and waveguides. A few prior works have explored the impact of PV on DWDM-based photonic links at the system-level [9][24]. In [9], a thermal tuning based approach is presented that adjusts chip temperature using dynamic voltage and frequency scaling (DVFS) to compensate for chip-wide PV-induced resonance shifts in MRs. In [24], a methodology to salvage network-bandwidth loss due to PV-drifts is proposed, which reorders MRs and trims them to nearby wavelengths. *All of these PV-remedial techniques are network specific and ignore the harmful effects of PV remedies on crosstalk.* Our proposed framework in this paper is different and novel as it considers the deleterious effects of IM crosstalk and PV-remedial techniques that increase crosstalk noise in detector MRs.




Fig. 1: Impact of PV-induced resonance shifts on MR operation in DWDM-based waveguides (note: only PV-induced red resonance shifts are shown): (a) MR as active modulator modulating in resonance wave-length with PV-induced red shifts (b) MR as active detector detecting its resonance wavelength with PV-induced red shifts.

3. PV-AWARE CROSSTALK ANALYSIS

3.1 Impact of Localized Trimming on Crosstalk

An MR can be considered to be a circular photonic waveguide with a small diameter, not to be confused with the larger DWDM-based photonic waveguide for which MRs serve as modulators and detectors. Variations in MR dimensions due to PV cause a “shift” in the resonance wavelengths of MRs. Fig. 1 shows the impact of PV on crosstalk noise (as dotted/dashed lines) in MRs. From Fig. 1(a) it can be seen that PV-induced red shifts in MR modulators increase crosstalk noise in the waveguide and decrease signal strength of non-resonating wavelengths. Fig. 1(b) shows how PV-induced red shifts increase detected crosstalk noise and decrease detected signal power of resonance wavelengths in MR detectors, which in turn reduces SNR and photonic data communication reliability.

As discussed earlier, the localized trimming method is essential to deal with PV-induced resonance red shifts in MRs. However, the use of this method in an MR alters its intrinsic optical properties, which leads to increased crosstalk noise and degraded performance in PNoCs that use these MRs. In this section, we discuss the effects of the localized trimming method on crosstalk and present analytical models to capture these effects in MRs. Further, we extend these models to generate system-level models for the Corona PNoC [1] in order to quantify signal and noise powers in the constituent MRs and DWDM waveguides of the Corona PNoC architecture.

The localized trimming method injects extra free carriers in the circular MR waveguide to counteract the PV-induced resonance red shifts. The introduction of extra free carriers reduces the refractive index of the circular MR waveguide, which in turn induces a blue shift in resonance to counteract the PV-induced red shifts. However, the extra free carriers increase the absorption related optical loss in the MR due to the free carrier absorption effect (FCA) [22]. The increase in the optical loss results in a decrease of MR Q-factor, which increases MR insertion loss and crosstalk, as discussed in Section 1.

We use a PV map (described in more detail in Section 4) to estimate PV-induced shifts in the resonance wavelengths of all the MRs across a chip. Then, for each MR device, we calculate the amount of change in refractive index (Δn_{si}) required to counteract this PV-induced wavelength shift using the following equation [21]:

$$\Delta n_{si} = \frac{\Delta \lambda_r * n_g}{\Gamma * \lambda_r}, \quad (1)$$

where, $\Delta \lambda_r$ is the PV-induced resonance shift that needs to be compensated for, λ_r is the target resonance wavelength of the MR, n_g is the group refractive index (ratio of speed of light to group velocity of all wavelengths traversing the waveguide) of the MR waveguide, and Γ is the confinement factor describing the overlap of the optical mode with the MR waveguide’s silicon core. We assume that the MR waveguides used in this study are similar to those reported in [22], fabricated using standard Si-SiO₂ material with a cross section of 450nm×250nm. The values of Γ and n_g for these MR waveguides are set to 0.7 and 4.2 respectively [22].

The required change in the free carrier concentration to induce the refractive index change of Δn_{si} at around 1.55 μ m wavelength

can be quantified using the following equation [21]:

$$\Delta n_{si} = -8.8 \times 10^{-22} \Delta N_e - 8.5 \times 10^{-18} (\Delta N_h)^{0.8}, \quad (2)$$

where, ΔN_e and ΔN_h are the change in free electron concentration and the change in free hole concentration respectively. The change in the absorption loss coefficient ($\Delta \alpha_{si}$) due to the change in free carrier concentration (owing to the FCA effect) can be quantified using the following equation [21]:

$$\Delta \alpha_{si} = -8.5 \times 10^{-18} \Delta N_e - 6.0 \times 10^{-18} \Delta N_h, \quad (3)$$

The Q-factor of an MR depends on this absorption loss coefficient. The relation between the Q-factor and $\Delta \alpha_{si}$, assuming critical coupling of MRs, is given by the following equation [22], where Q' is the loaded Q-factor of the MR:

$$Q' = Q + \Delta Q = \frac{\pi n_g}{\lambda_r(\alpha + \Delta \alpha_{si})}, \quad (4)$$

where, ΔQ is the change in Q-factor and α is the original loss coefficient, which is a sum of three components: (i) intrinsic loss coefficient due to material loss and surface roughness; (ii) bending loss coefficient, which is a result of the curvature in the MR; and (iii) the absorption effect factor that depends on the original free carrier concentration in the waveguide core. Typically, the localized trimming method injects excess concentration of free carriers into the MR, which increases the absorption loss coefficient (positive $\Delta \alpha_{si}$). As evident from Eq. (4), a positive value of $\Delta \alpha_{si}$ results in a decrease of the Q-factor. This causes a broadening of the MR passband, which results in increased insertion loss and crosstalk power penalties.

We model the MR transmission spectrum using a Lorentzian function. This function is used to represent coupling factor Φ in Eq. (5) between wavelength λ_i and an MR with resonance wavelength λ_j . From Eq. (3) and (4), it can be inferred that an MR's loaded Q-factor (Q') decreases with localized trimming. This in turn increases Φ and crosstalk noise. Further, using the same function, we determined loss factor γ in Eq. (6) which is the factor by which signal power of a wavelength λ_i is reduced when it passes through an MR whose resonance wavelength is λ_j . Through loss of a wavelength in a waveguide when it passes through an MR is defined as γ times the signal power of all wavelengths received before the MR.

$$\Phi(\lambda_i, \lambda_j, Q') = \left(1 + \left(\frac{2Q'(\lambda_i - \lambda_j)}{\lambda_j}\right)^2\right)^{-1}, \quad (5)$$

$$\gamma(\lambda_i, \lambda_j, Q') = \left(1 + \left(\frac{2Q'(\lambda_i - \lambda_j)}{\lambda_j}\right)^2\right)^{-1}, \quad (6)$$

In the next section, we use the derived values of coupling factor Φ and loss factor γ from this section to model worst case crosstalk and SNR for the Corona PNoC, in the presence of process variations.

3.2 PV-Aware Crosstalk Models for Corona PNoC

We characterize crosstalk in DWDM-waveguides for the well-known Corona PNoC enhanced with token-slot arbitration [1], [10]. In DWDM-based waveguides, data transmission requires modulating light using a group of MR modulators equal to the number of wavelengths supported by DWDM. Similarly, data detection at the receiver requires a group of detector MRs equal to the number of DWDM wavelengths. We present analytical equations to model worst-case crosstalk noise power, maximum power loss, and SNR in detector MR groups (similar equations are applicable to modulator MR groups). Before presenting actual equations, we provide notations for the parameters used in the equations, in Tables I and II.

The Corona PNoC is designed for a 256 core single-chip platform, where cores are grouped into 64 clusters, with 4 cores in each cluster. A photonic crossbar topology with 64 data channels is used for communication between clusters. Each channel consists of 4 multiple-write-single-read (MWSR) waveguides with 64-wavelength DWDM in each waveguide. As modulation occurs on both positive and negative edges of the clock in Corona, 512 bits (cache-line size) can be modulated and inserted on 4 MWSR waveguides in a single cycle by a sender. A data channel starts at a cluster called 'home-cluster', traverses other clusters (where modulators can modulate

light and detectors can detect this light), and finally ends at the home-cluster again, at a set of detectors (optical termination). A power waveguide supplies optical power from an off-chip laser to each of the 64 data channels at its home-cluster, through a series of 1X2 splitters. In each of the 64 home-clusters, optical power is distributed among 4 MWSR waveguides equally using a 1X4 splitter with splitting factor R_{S14} . As all 1X2 splitters are present before the last (64th) channel, this channel suffers the highest signal power loss. Thus, the worst-case signal and crosstalk noise exists in the detector group of the 64th cluster node, and this node is defined as the worst-case power loss node (N_{WCPL}) in the Corona PNoC.

Table I: Notations for photonic power loss, crosstalk coefficients [15]

Notation	Parameter type	Parameter value (in dB)
L_p	Propagation loss	0.274 per cm
L_b	Bending loss	-0.005 per 90°
L_{S12}	1X2 splitter power loss	-0.2
L_{S14}	1X4 splitter power loss	-0.2
L_{S16}	1X6 splitter power loss	-0.2

Table II: Other model parameter notations

Notation	Crosstalk Coefficient	Parameter Value
Q	Q-factor	9000
L	Photonic path length in cm	
B	Number of bends in photonic path	
λ_j	Resonance wavelength of MR	
R_{S12}	Splitting factor for 1X2 splitter	
R_{S14}	Splitting factor for 1X4 splitter	
R_{S16}	Splitting factor for 1X6 splitter	

For this N_{WCPL} node, the signal power ($P_{\text{signal}}(\lambda_j)$) and crosstalk noise power ($P_{\text{noise}}(\lambda_j)$) received at each detector with resonance wavelength λ_j are expressed in Eq. (7) and (8). $P_s(\lambda_i, \lambda_j)$ in Eq. (9) is the signal power of the λ_i wavelength received before the detector with resonance wavelength λ_j . $K(\lambda_i)$ in Eq. (11) represents signal power loss of λ_i before the detector group of N_{WCPL} . $\psi(\lambda_i, \lambda_j)$ in Eq. (10) represents signal power loss of λ_i before the detector with resonance wavelength λ_j within the detector group of N_{WCPL} . Due to PV, crosstalk coupling factor (Φ , Eq. (5)) increases with decrease in loaded Q-factor (Q', Eq. (4)), which in turn increases crosstalk noise in the detectors. We can define SNR(λ_j) of the detector having resonance wavelength λ_j of N_{WCPL} as the ratio of $P_{\text{signal}}(\lambda_j)$ to $P_{\text{noise}}(\lambda_j)$, as shown in Eq. (12).

$$P_{\text{signal}}(\lambda_j) = \Phi(\lambda_j, \lambda_j, Q'_{(63 \times 64) + j}) P_s(\lambda_j, \lambda_j), \quad (7)$$

$$P_{\text{noise}}(\lambda_j) = \sum_{i=1}^n \Phi(\lambda_i, \lambda_j, Q'_{(63 \times 64) + j}) (P_s(\lambda_i, \lambda_j)) (i \neq j), \quad (8)$$

$$P_s(\lambda_i, \lambda_j) = K(\lambda_i) \psi(\lambda_i, \lambda_j) P_{\text{in}}(i), \quad (9)$$

$$\psi(\lambda_i, \lambda_j) = \prod_{k=1}^{(k-1) < j} \gamma(\lambda_i, \lambda_k, Q'_{(63 \times 64) + k}), \quad (10)$$

$$K(\lambda_i) = (R_{S14})(L_{S14})(L_p)^L (L_b)^B \prod_{n=1}^{63} \prod_{j=1}^{64} \gamma(\lambda_i, \lambda_j, Q'_{((n-1) \times 64) + j}), \quad (11)$$

$$\text{SNR}(\lambda_j) = \frac{P_{\text{signal}}(\lambda_j)}{P_{\text{noise}}(\lambda_j)}, \quad (12)$$

3.3. Modeling PV of MR Devices in Corona PNoC

We adapt the VARIUS tool [23] similar to prior work [24] to model die-to-die (D2D) as well as within-die (WID) process variations in MRs. We consider photonic devices with a silicon (Si) core and silicon-dioxide (SiO₂) cladding. VARIUS uses a normal distribution to characterize on-chip D2D and WID process variations. The key parameters are mean (μ), variance (σ^2), and density (α) of a variable that follows the normal distribution. As wavelength variations are approximately linear to dimension variations of MRs, we assume they follow the same distribution. The mean (μ) of wavelength variation of an MR is its nominal resonance wavelength. We consider a DWDM wavelength range in the C and L bands [14], with a starting

wavelength of 1550nm and a channel spacing of 0.8nm. Hence, those wavelengths are the means for each MR modeled. The variance (σ^2) of wavelength variation is determined based on laboratory fabrication data [6] and our target die size. We consider a 256-core chip with die size 400 mm² at a 22nm process node. For this die size we consider a WID standard deviation (σ_{WID}) of 0.61nm [24] and D2D standard deviation (σ_{D2D}) of 1.01 nm [24]. We also consider a density (α) of 0.5 [24] for this die size. With these parameters, we use VARIUS to generate 100 process variation maps. Each process variation map contains over one million points indicating the PV-induced resonance shift of MRs. The total number of points picked from these maps equal the number of MRs in the Corona PNoC.




Fig. 2: Transmission spectrum of MR groups with (a) high channel gap (CG) (b) low channel gap (CG); (C) IMCM at low channel gap.

4. IM CROSSTALK ANALYSIS

Intermodulation (IM) crosstalk occurs when a resonance wavelength of an MR modulator is modulated by the neighboring MR modulators. As evident from Eq. (6), signal strength of wavelengths in photonic waveguides of DWDM-based PNoCs decrease with increase in loss factor (γ). This γ increases with a decrease in IM gap, which is the gap between resonance wavelengths of an MR in active and inactive state. Furthermore, this reduction in signal strength of the resonance wavelength also depends on the channel gap (CG) between two adjacent wavelengths in the DWDM. Fig. 2 shows the transmission spectrum of MR groups with high and low CG. A change from low DWDM (Fig. 2(a)) to higher DWDM (Fig. 2(b)) reduces the CG and IM gap, which in turn increases IM crosstalk as is evident from the intersection of the transmission spectrum of inactive MRs with wavelengths in the waveguide (λ_1 - λ_n). This IM crosstalk increases wavelength signal loss.

All MR modulators at a node in a DWDM waveguide have neighbors on both sides except the first and the last modulators. So, the first (λ_1) and the last (λ_n) wavelengths of DWDM have the lowest signal losses and highest signal strengths. Thus, the modulated set of DWDM wavelengths that travel along a photonic waveguide to the target detector node have varying signal strengths. At an MR detector group, the first (λ_1) wavelength signal gets filtered and detected by the first detector. As a result, the signal strength of the first wavelength becomes negligible. This negligible signal strength of the first wavelength does not significantly add crosstalk noise in the succeeding neighboring detectors. In contrast, the last (λ_n) wavelength, which also has higher signal strength, gets filtered and detected by the last detector in the detector group. So, the last (λ_n) wavelength signal has to travel along all the detectors in the group of detector rings before being detected. On its way to the last detector, the last wavelength signal incurs crosstalk noise in all the detectors across

the detector group. As the strength of the last (λ_n) wavelength signal is high, the incurred crosstalk noise is also high.

5. IM-AWARE CROSSTALK MITIGATION

Based on the observations in the previous section, we propose an IM passband truncation aware crosstalk mitigation (IMCM) scheme to decrease crosstalk noise in MRs of DWDM based photonic links. In IMCM, to reduce signal strength of the last wavelength in the DWDM, we propose placing an additional MR at each modulating and detecting node. This extra MR is tuned near to the last (λ_n) wavelength of DWDM with a tuning distance of half the channel gap (CG/2) of the DWDM (as shown in Fig. 2(c)). This extra MR increases signal loss of this last (λ_n) wavelength and reduces its signal strength. Thus, it creates uniform signal loss across all wavelengths used in the DWDM. This extra MR (passband of this MR is shown with a dotted line in Fig. 2(c)) is always maintained in inactive mode and reduces the effects of IM crosstalk on the boundary wavelengths of DWDM by reducing their respective signal strengths. This mechanism reduces crosstalk in detecting MRs to improve SNR (and thus reduce BER). To implement the IMCM technique in the Corona PNoC, there is a need to increase the number of MRs in all modulating and detecting nodes by one on their MWSR and SWMR waveguides. The increase in MRs on the waveguides increases throughput and laser power. We account for this overhead in our analysis.




Fig. 3: Overview of proposed PVCM technique

6. PV-AWARE CROSSTALK MITIGATION

We also propose a PV-aware trimming-induced crosstalk mitigation (PVCM) scheme, which is illustrated in Fig. 3. PV-induced red shifts can be realigned using localized trimming, but this process worsens crosstalk noise. From Eq. (5), crosstalk in MR detectors of DWDM-based PNoCs increases with increase in coupling factor (Φ) and increase in signal strength of an immediate non-resonating wavelength. This implies that the trimming-affected crosstalk in a detector can be reduced by reducing the signal strength of immediate non-resonating wavelengths. Therefore, our proposed PVCM technique decreases the signal strength of the immediate non-resonant wavelength by modulating a zero (shielding bit) on it, which reduces crosstalk noise in the detector. The PVCM technique first divides detecting MRs into groups of 8 MRs each. Then, it determines the maximum PV-induced resonance red shift ($\Delta\lambda_{max}$) in each MR group. As discussed in [17], the PV-induced resonance shifts in MRs can be gauged in situ at system initialization by using a dithering signal to generate an anti-symmetric error signal that indicates the magnitude of PV-induced resonance shifts. The overhead of this in-situ PV detection technique can be considered to be negligible [17]. In our analysis, we model and estimate PV in MRs using the VARIUS tool [23], a description of which was given in Section 3.3.

Once PV-induced red shifts of MRs are determined, we store information about whether to enable or disable encoding (i.e., injecting shield bits between data bits) for each MR group in a read-only memory (ROM) at the modulating node, based on the maximum PV-induced resonance red shift ($\Delta\lambda_{max}$) value for the group. If this value is greater than a threshold red shift value ($\Delta\lambda_{th}$) for an MR group, we store a '1' to enable PVCM, else we store a '0' to disable PVCM for this MR group. MR groups with $\Delta\lambda_{max} < \Delta\lambda_{th}$ are thus not impacted. Only MR-groups with $\Delta\lambda_{max} > \Delta\lambda_{th}$ employ encoding.

7. PICO FRAMEWORK: SENSITIVITY ANALYSIS

We combine the IMCM scheme that mitigates the effects of IM crosstalk and the PVCMM scheme that mitigates the PV-affected crosstalk in PNoCs into a holistic crosstalk mitigation framework called *PICO*. As the number of shield bits used in *PICO* increases, laser power and trimming power of PNoCs also increase. Thus, we need to limit the number of shield bits. We performed a sensitivity analysis using the Corona PNoC with varying number of shield bits per detector node to quantify its effect on worst-case SNR. We analyzed worst case SNR with 0%, 25%, 50%, 75% and 100% of shield bits added to data bits for the Corona PNoC. Based on our analysis across 100 process variation maps, we determined the value of $\Delta\lambda_{th}$ to be 0.45nm, 0.88 nm, 1.25nm and 4.25nm, for the cases with 25%, 50%, 75% and 100% of shielding bits to data bits, respectively.




Fig. 4: Sensitivity analysis in terms of worst-case SNR for Corona PNoC with *PICO* allowing 0%, 25%, 50% and 100% ratio of shield bits to data bits across 100 process variation maps; average power consumption for each configuration is also shown on the top of each bar.

Fig. 4 shows the range of worst-case SNR values across PV maps, for different ratios of shield bits to data bits. From the figure it can be seen that on average *PICO* with 25%, 50%, 75% and 100% shield bits has 8.2%, 19.77%, 26.5% and 40.9% higher worst-case SNR (note: higher SNR is better) respectively compared to the baseline (with 0% shielding). Intuitively, higher ratios of shield bits should result in higher worst case SNR, as more shield bits can be used to protect data bits, which in turn reduces crosstalk and improves SNR. But, with increase in number of shield bits, the number of MRs on the waveguides increases, which increases the through losses, requiring more laser power to compensate for the losses. Addressing PV drifts for high MR counts also requires higher trimming power in PNoCs. Fig. 4 shows that average power consumption with 25%, 50%, 75% and 100% shield bits is 12.6%, 33.5%, 62.2% and 109.5% higher compared to the baseline. To balance crosstalk reliability and power overheads, we select the 50% shield bits to data bits configuration for the rest of our experiments.

To implement our *PICO* framework with 50% shielding bits on the Corona PNoC, we increase the number of MWSR waveguides in each channel from 4 to 6, to maintain the same bandwidth as in the baseline case. Additionally each modulating node needs to store 2,646 bits in its ROM to capture encoding requirements for all the remaining 63 detecting nodes. Power and area overheads for these modifications are presented in the next section. Lastly, we also consider up to a two cycle overhead for encoding and decoding of data in *PICO*. The first cycle is needed to retrieve data from the ROM storage, whereas the second cycle is used only if data is to be encoded before sending on the waveguide.

8. EXPERIMENTS

8.1. Experimental Setup

To evaluate our proposed crosstalk noise mitigation framework *PICO* (IMCM+PVCMM) in DWDM-based PNoCs, we implement and integrate it with the Corona [1], [10] crossbar-based PNoC. We modeled and performed simulation based analysis of the enhanced Corona PNoC using a cycle-accurate NoC simulator, for a 256 core single-chip architecture at 22nm. As explained in Section 3.3, we generated 100 PV maps to evaluate how *PICO* performs for different PV profiles. We used real-world traffic from applications in the PARSEC benchmark suite [11]. GEM5 full-system simulation [16]

of parallelized PARSEC applications was used to generate traces that were fed into our cycle-accurate NoC simulator. We set a “warm-up” period of 100 million instructions and then captured traces for the subsequent 1 billion instructions. We performed geometric calculations for a 20mm×20mm chip size, to determine lengths of MWSR waveguides in the Corona PNoC. Based on this analysis, we estimated the time needed for light to travel from the first to the last node as 8 cycles at 5 GHz clock frequency. We use a 512 bit packet size, as advocated in the Corona PNoC.

The static and dynamic energy consumption of electrical routers and concentrators in Corona is based on results from the open source DSENT tool. We model and consider area, power, and performance overheads for our framework implemented with the Corona PNoC, as follows. *PICO* has an electrical area overhead estimated to be 6.24 mm² and a power overhead of 1.14 W, using gate-level analysis and the CACTI 6.5 [18] tool for memory and buffers. The photonic area overhead is 9.44 mm², based on the physical dimensions [14] of waveguides, MRs, and splitters. For energy consumption of photonic devices, we adapt model parameters from recent work [15], [19], [20], with 0.42pJ/bit for every modulation and detection event and 0.18pJ/bit for the driver circuits of modulators and photodetectors. We used optical loss for photonic components, as shown in Table I, to determine the photonic laser power budget and correspondingly the electrical laser power. The MR trimming power is set to 130μW/nm [25] for current injection (blue shift).

8.2. Experimental Results with Corona PNoC

Our first set of experiments compares the baseline Corona PNoC with fair token-slot arbitration [1], [10] but without any crosstalk-enhancements, with three variants of the architecture corresponding to the three crosstalk-mitigation strategies we compare: PCTM5B and PCTM6B from [12] and our proposed *PICO* framework from this paper. PCTM5B and PCTM6B are encoding schemes that replace each 4-bits of a data word with 5-bit and 6-bit code words. These schemes aim to reduce photonic signal-strength of immediate non-resonant wavelengths (adjacent wavelengths in DWDM) to decrease crosstalk and improve SNR in MR detectors.




Fig. 5: Worst-case SNR comparison of *PICO* with PCTM5B [12] and PCTM6B [12] for Corona PNoC considering 100 process variation map

Utilizing the models presented in section 3, we calculate the received crosstalk noise and SNR at detectors for the node with worst-case power loss (N_{WCPL}), which corresponds to MR detectors in cluster 64 for the Corona PNoC. While the worst-case SNR for the baseline Corona PNoC occurs when all of the 64-bits of a received data word in a waveguide are 1’s, for the implementations of Corona with PCTM5B, PCTM6B and *PICO*, this is not the case, i.e., each detector in cluster 64 has a worst-case SNR for a different pattern of 1’s and 0’s in the received data word. We used our analytical models to determine these unique worst-case patterns for each of the techniques when used with Corona, for an accurate analysis.

Fig. 5 summarizes the worst-case SNR results for the baseline, PCTM5B, PCTM6B, and *PICO*. From the figure, it can be observed that Corona PNoC with *PICO* has 4.4×, 2.05×, and 1.2× SNR improvements on average, compared to baseline, PCTM5B, and PCTM6B respectively. Both the PCTM5B and PCTM6B techniques eliminate occurrences of ‘111’ in a data word and have limited occurrences of ‘11’, which helps reduce crosstalk noise in the detectors. But these techniques do not consider the impact of IM effects

and PV resonance wavelength drifts. More specifically, IM can create significant additional crosstalk with these techniques in some cases where occurrences of ‘11’ are present. PV in MRs also varies signal power of wavelengths in DWDM as they propagate through the waveguide, so there is need for encoding on specific wavelengths where there is high signal loss (due to trimming) which is not considered in both PCTM5B and PCTM6B. Due to these reasons, PCTM5B and PCTM6B have worse SNR degradation. *PICO* reduces crosstalk in the detectors by combining benefits from IMCM and using PVCMM’s shield bits between data bits. *PICO* also considers the PV profile of MRs to intelligently select MRs for shielding.




Fig. 6: (a) normalized latency and (b) energy-delay product (EDP) comparison between Corona baseline and Corona with PCTM5B, PCTM6B, and *PICO* techniques, for PARSEC benchmarks. Latency results are normalized to the baseline Corona architecture results.

Fig. 6 (a) and (b) present detailed simulation results that quantify the average network packet latency and energy-delay product (EDP) for the four Corona configurations. Results are shown for twelve multi-threaded PARSEC benchmarks. From Fig. 6(a) it can be seen that on average, Corona with *PICO* has 12.6% higher latency compared to baseline, and it also has 2.1% higher latency compared to both PCTM5B and PCTM6B. The additional delay due to encoding and decoding of data with *PICO*, PCTM5B and PCTM6B contributes to their increase in average latency. The penalty due to encoding/decoding is 1 cycle in PCTM5B and PCTM6B, whereas *PICO* has a 1 or 2 cycle penalty, which increases its delay overhead.

From the results for EDP shown in Fig. 6(b), it can be seen that on average, the Corona configuration with our *PICO* framework has 17.2% higher EDP compared to the baseline. Increase in EDP for Corona with *PICO* is not only due to the increase in average latency, but also due to the addition of extra bits for encoding and decoding, which leads to an increase in the amount of photonic hardware in the architectures (more number of MRs, complex splitters). This in turn increases static energy consumption. Dynamic energy also increases in these architectures, but by much less. However, EDP for the *PICO* framework is 5.1% and 16.18% lower compared to PCTM5B and PCTM6B respectively. Despite the higher latency overhead compared to PCTM5B, *PICO* saves considerable dynamic energy compared to PCTM5B as it uses lower number of bits for traversal of the packet. In a similar manner, although *PICO* has higher latency compared to PCTM6B, it conserves laser and trimming/tuning power due to lower photonic hardware requirements than PCTM6B.

9. CONCLUSION

We have presented a novel heterodyne crosstalk mitigation framework for the reduction of crosstalk noise in the detectors of DWDM based photonic network-on-chip (PNoC) architectures. Our proposed *PICO* framework shows interesting trade-offs between reliability, performance, and energy overhead for the Corona crossbar-based PNoC architecture. Our experimental analysis shows that the *PICO* framework improves worst-case SNR by 4.4 \times compared to the baseline Corona PNoC architecture, and by up to 2.05 \times compared to the best known PNoC crosstalk mitigation schemes from prior work. Thus, *PICO* represents an attractive solution to enhance reliability in emerging DWDM-based PNoCs.

10. ACKNOWLEDGMENTS

This research is supported by grants from SRC, NSF (CCF-1252500, CCF-1302693), and AFOSR (FA9550-13-1-0110).

REFERENCES

- [1] D. Vantrease et al., “Corona: System implications of emerging nanophotonic technology,” in Proc. ISCA, 2008.
- [2] Y. Pan et al., “Firefly: Illuminating future network-on-chip with nanophotonics,” in Proc. ISCA, 2009.
- [3] S. V. R. Chittamuru et al., “A Reconfigurable Silicon-Photonic Network with Improved Channel Sharing for Multicore Architectures,” in ACM GLSVLSI, May, 2015.
- [4] L.H.K. Duong et al., “Coherent Crosstalk Noise Analyses in Ring-based Optical Interconnects,” in Proc. DATE, 2015.
- [5] K. Padmaraju et al., “Intermodulation Crosstalk Characteristics of WDM Silicon Microring Modulators,” in PTL, 2014.
- [6] S. K. Selvaraja., “Wafer-Scale Fabrication Technology for Silicon Photonic Integrated Circuits,” PhD thesis, Ghent University, 2011.
- [7] S. V. R. Chittamuru et al., “Improving Crosstalk Resilience with Wavelength Spacing in Photonic Crossbar-based Network-on-Chip Architectures,” IEEE MWSCAS, Aug. 2015.
- [8] C. Batten et al., “Building manycore processor-to-dram networks with monolithic silicon photonics,” in HotI, pages 21–30, 2008.
- [9] Z. Li et al., “Reliability modeling and management of nanophotonic on-chip networks,” IEEE TVLSIS, 20:98–111, 2010.
- [10] D. Vantrease et al., “Light speed arbitration and flow control for nanophotonic interconnects,” in Proc. IEEE/ACM MICRO, 2009.
- [11] C. Bienia et al., “The PARSEC Benchmark Suit: Characterization and Architectural Implications,” in PACT, Oct. 2008.
- [12] S. V. R. Chittamuru et al., “Crosstalk Mitigation for High-Radix and Low-Diameter Photonic NoC Architectures,” in IEEE D&T, 2015.
- [13] Q. Xu, B. Schmidt et al., “Cascaded silicon micro-ring modulators for wdm optical interconnection,” in Opt. Exp., vol. 14, 2006.
- [14] S. Xiao et al., “Modeling and measurement of losses in silicon-organic resonators and bends,” in Opt. Exp., vol.15, no.17, 2007.
- [15] L.H.K. Duong et al., “A Case Study of Signal-to-Noise Ratio in Ring-Based Optical Networks-on-Chip,” in IEEE D&T, 2014.
- [16] N. Binkert et al., “The gem5 Simulator,” in CA News, May 2011.
- [17] K. Padmaraju and K. Bergman, “Resolving the thermal challenges for silicon microring resonator devices,” Nanophotonics, 2 (4), 2013.
- [18] CACTI 6.5, <http://www.hpl.hp.com/research/cacti/>
- [19] X. Zheng et al., “Ultra-efficient 10Gb/s hybrid integrated silicon photonic transmitter and receiver,” in Opt. Express, Mar 2011.
- [20] P. Grani and S. Bartolini, “Design Options for Optical Ring Interconnect in Future Client Devices,” in ACM JETC, May, 2014.
- [21] R. G. Beausoleil, “Large-Scale Integrated Photonics for High-Performance Interconnects,” ACM JETC, Vol. 7, No. 2, 2011.
- [22] K. Preston, et al., “Performance guidelines for WDM interconnects based on silicon microring resonators,” IEEE CLEO, 2011.
- [23] S. Sarangi et al., “Varius: A model of process variation and resulting timing errors for microarchitects,” IEEE TSM, 21(1):3–13, 2008.
- [24] Y. Xu et al., “Tolerating process variations in nanophotonic on-chip networks,” in Proc. ISCA, Portland, OR, USA, 2012, pp. 142–152.
- [25] C. Nitta et al., “Addressing system-level trimming issues in on-chip nanophotonic networks,” in Proc. HPCA, 2011.